Why Use a 99¢ Reserve Price on eBay?

Joseph Uri Podwol
Department of Economics, Cornell University

July 8, 2006

Abstract

The popularity of the 99¢ reserve price in auctions held on the eBay platform provides an interesting economic puzzle. That is, the seller's receipts are maximized using a 99¢ reserve price for many items. This paper presents a model showing that when the seller cannot commit to a reserve price, the best she can do is set the reserve low enough so as to be non-binding and allow the price to be determined by the bidding.

1 Introduction

The growing literature on online auctions has generated several new ideas aiding our understanding of how traditional markets work. This paper takes the opposite approach, using an old idea to explain an interesting development in the workings of online markets, in particular the popularity of the 99¢ reserve price in eBay auctions. The 99¢ reserve is popular for one simple reason: it works. The reason it works is something of a puzzle to eBay merchants. The following excerpt, taken from an eBay message board, describes the nature of the puzzle:

> I have been in the practice of listing items for the least amount of money I would be willing to sell them for. I figure my costs and labor and go from there.
>
> I discovered that another seller was selling items similar to mine and starting them at 99 cents. I thought that this seller must be nuts to sell stuff for 99 cents. I noticed that many of her bids went way above what my similar items started at, where many of my items started at the higher price just stagnated ending without bids.
>
> I have been experimenting with this 99 cent auction thing I have noticed that the bids do much better and go higher than what I normally would have started them at. Sure, some sell for 99 cents but the others that go higher balance it all out[.]
>
> –Treasures_by_cynthia

The puzzle, as identified by this retailer, is the following: Two identical items are put up for auction at different times using different posted reserve prices, say 99¢ and $14 respectively. The sales mechanism is a second-price auction in both cases.[1] The auction with the 99¢ reserve ends with a price of $16 while the $14 reserve auction fails to generate any bids. This outcome is a puzzle to this and other merchants as a bidder willing to pay $16 in a 99¢ reserve auction should be willing to do so in a $14 reserve auction. Another retailer, after conducting a number of experiments on the issue, advises that the best use of the 99¢ auction is only when 99¢ is a small fraction of the item's market value:

> 99 cent auctions are best for items with values (completed auction, or perceived value) of over $10.00, even better for items worth over $20.00. It doesn't work as well with items worth 5-6 dollars.
> –Clact

Explanations for this puzzle abound on the message board. The most well-thought-out rely on what could be deemed "behavioral" hypotheses. One merchant claims that once the bidding begins, a consumer's objective changes from maximizing expected surplus to winning the auction with the highest probability. Another identifies a "quasi-endowment" effect that does not allow a buyer to let go of the object after being outbid by a competitor.

The logic of these behavioral hypotheses notwithstanding, the merchants appear to be overlooking the dynamic effect inherent in what appears to be a common practice. That is, when the item fails to sell at the $14 reserve price, they simply "re-list" with a lower reserve. Anticipating this response, rational consumers may wish to hold off bidding in the current auction and instead wait for the next auction in which they may obtain the item at a lower price. For instance, a potential buyer who values the item at $14 gains nothing by winning the current auction at a price of $14, but earns positive surplus in the following period should the item fail to sell presently. The seller, anticipating this response, is then left with no choice but to start at the lower reserve. It is this wait-and-see approach by buyers that accounts for the puzzle identified by our seller Treasures_by_cynthia. Since the item will necessarily sell when the reserve is 99¢, the buyer with a valuation of $14 can do no better than bid his valuation in the current period.

This dynamic is essentially a problem of commitment identified in the bargaining literature by Schelling (1960) and in the durable goods monopoly literature by Coase (1972). In Coase's model, a monopolist owns a tract of land to be sold off in lots. Charging a supra-competitive price to the batch of customers with the highest willingness to pay, there are lots left over to be sold to the next at a lower price. This process continues until the last lot is sold at marginal cost

[1] Buyers are given a predetermined window, say three days, during which they may place bids. As bids are received, the current price as well as the high bidder's identity are made public. At the end of the auction, the high bidder wins the item and pays an amount equal to the second-highest bid plus a small increment. Given the ending rule, the second-price auction model seems most appropriate discription of the environment. For a more complete description of the auction environment, see Bajari & Hortacsu (2003).

and in the absence of frictions the price reaches marginal cost “in the twinkling of an eye.”(2) In equilibrium, no customer will pay more than the competitive price for a lot of land. In the case of bargaining, a mutually beneficial transaction can be precluded by a problem of indeterminacy, as any division of surplus is equally plausible. However, when one of the parties can make a *commitment* preventing him from accepting any unfavorable offers, "he can squeeze the range of indeterminacy down to the point most favorable to him."(13). Likewise, if the durable goods monopolist could commit himself against future price reductions, he can regain his power to price as a monopolist.

This purpose of this paper is to show that the seller’s inability to commit to a reserve price can account for the popularity of the 99¢ reserve. In doing so, we show that this model explains two empirical anomalies: 1. Bidding gaps: An auction with a high reserve, say $14, ends with no bids while an auction with a non-binding reserve of 99¢ ends with a price exceeding $14. Empirically, this means that for any price of p, the probability that the sales price exceeds p is decreasing in the reserve. 2. Reserve optimality: A non-binding reserve of 99¢ maximizes the seller’s expected revenue.

The optimality of a non-binding reserve price has been documented in the empirical literature. Reiley (2004) conducts field experiments auctioning Magic cards over the internet via first-price sealed-bid auction and finds that a zero reserve maximizes expected revenues. Bajari & Hortacsu (2003) find that sellers of rare coins on eBay tend to post reserve prices well below the coin’s book value. A possible explanation given by Reiley, owing to the model of Levin & Smith (1994), is that participation is endogenous and that buyers face participation costs. In such a world, a buyer is less likely to participate in an auction with a higher reserve as his expected surplus may not be large enough to compensate for having paid the cost of participation. On average, a low-reserve high-expected-participation auction yields greater revenue than a high-reserve low-expected-participation auction. Riley claims to find evidence of participation costs by auctioning several items simultaneously with no reserve and finds a paucity of bidders placing bids on all items.

Another possible explanation is that buyers’ valuations are affiliated. In the affiliated model of Milgrom & Weber (1982), not only are the valuations of bidders interdependent, but information held by one buyer is positively correlated with the valuation of another. In this way, a bid placed by one bidder may cause another to reevaluate his willingness to pay and increase his bid. Thus, the greater the number of bids, the more opportunities for other bidders to update their bids to the benefit of the seller. Thus the seller is well served to start the auction at a low reserve.

With the existing theories in mind, this paper offers an alternative model to explain the prevalence of the 99¢ reserve. This model distinguishes itself from the explanations given on the message board in that it does not rely on irrational behavior. And contrary to the model of Milgrom & Weber, it is applicable to markets in which buyers valuations would not be expected to be interdependent. And as we will see, the model of Levin & Smith (1994), while plausible within this context, is ambiguous as to the presence of bidding gaps.

Section 2 presents both the standard auction model and the dynamic model in which the seller cannot commit. Section 3 addresses the economic implications of the dynamic model to explain the aforementioned empirical anomalies in a way the standard model cannot. Section 4 contrasts the relevant models and derives empirical tests to distinguish one from the next, and Section 5 concludes.

2 Independent Private Values Model with Listing Fees

Consider an auction marketplace consisting of a single seller and n potential buyers indexed $i = 1, 2, ..., n$. The seller has one unit of a particular item to sell. Buyers are risk neutral, have unit demands, and differ only in their valuation of the item. Each buyer's valuation, denoted v, is private information. Valuations are assumed to be independent and identically distributed according to F, a continuous distribution with density f defined over the support $[\underline{v}, \bar{v}]$, $\underline{v} \geq 0$. With the cost of production sunk, the seller has no value for the item, is risk neutral, and offers the item for sale via sealed bid second-price auction with posted reserve $r \in [0, \infty)$. For the privilege of running the auction on the online platform, the seller pays a listing fee $c(r)$, weakly increasing in the posted reserve. The schedule of listing fees, consistent with what is used on eBay, $c(r)$ is a step function taking on values c^0 for $r \in [0, r^0]$, c^1 for $r \in (r^0, r^1]$ and so forth.[2] Assume the seller's valuation, the distribution of buyer valuations F and the schedule of listing fees $c(r)$ are all common knowledge. Lastly, we restrict the shape of F to guarantee concavity of the seller's objective function so that the seller's optimal reserve is unique.

This setup is meant to capture an environment of asymmetric information in which each buyer knows his willingness to pay but not that of other buyers. The seller has only imperfect information over the valuation of each buyer and must set an optimal reserve price given her information. Since their valuations are known to themselves, buyers learn nothing about the item's value from the seller's reserve price or from the bidding of others. Through this environment, we compare the optimal reserve price, the average resultant transaction price, and the seller's net revenue when: 1.commitment by the seller is possible and 2. commitment is not possible. By commitment, I am referring to a credible promise to not offer the item for sale at a future date with a lower reserve price.[3] Thus commitment could take the form of a promise to either burn the merchandise following the auction in the event the auction does not result in a sale, or to keep the reserve price at a previously determined level when the item goes up for auction again. Either promise, if credible, eliminates any incentive

[2] The schedule of listing fees actually used on eBay is given in Table 1 of the Appendix.

[3] This "commitment" is analogous to that referred to in the durable goods monopoly literature whereby the monopolist can prevent himself from lowering the sale price in future periods.

on the part of buyers to wait for the item to be listed at a lower price and buy it then.

In the literature on optimal auctions,[4] it is commonly assumed that a seller can commit to no future sales in the event the item goes unsold at the optimal reserve price, or any reserve price for that matter. While this assumption may seem valid when the seller is a reputable auction house, it is clearly not so when the seller is one of thousands of semi-anonymous eBay merchants. In various online formats, these merchants readily admit amongst themselves to the practice of relisting, seemingly unaware of the consequent dynamic effects. Further, the relative anonymity of retailers makes establishing a reputation difficult and pledges not to relist non-credible.

2.1 Standard Auction Model: Seller Commitment

For means of comparison, consider a one-shot game where the seller commits to a given reserve r. The auction game proceeds as follows: Buyers are called upon to submit sealed bids.[5] If the highest bid exceeds r, the item sells for the greater of the second-highest bid and r. If not, all players get nothing and the game ends. In the event of a tie, the item is awarded to one of the tying bidders at random and the winner pays his bid. The equilibrium concept is *Bayesian-Nash equilibrium.*

This game is identical to the second-price auction mechanism devised by Vickrey (1961) with the added feature of a reserve price. Vickrey showed that each bidder follows a weakly dominant strategy of revealing his true type v through his bid. The following lemma shows that Vickrey's strategy extends to an environment with a reserve price.

Lemma 1 *In the second-price auction with seller commitment and reserve r,the following bidding strategy is (weakly) dominant: Bid v if $v \geq r$; do not bid otherwise.(7)*

Proof. Consider the strategy of buyer 1 with valuation v and let b_1 denote the maximum of the bids placed by all other buyers. There are 3 cases to consider. In each case, we see that buyer 1's payoffs are at least as great playing the proposed strategy as they would be under any possible deviation.

1. $v > b_1$:If $v > r$ then deviating affects the outcome only when player 1 reduces his bid so as to go from winning the auction and paying $\max\{r, b_1\} < v$ to losing the auction and getting nothing. If $v < r$ then deviating affects the outcome only when player 1 bids an amount greater than r so as to go from receiving nothing by not bidding, to winning the auction, paying r, and receiving negative surplus.

[4]See Riley & Samuelson (1981) for an intuitive derivation of the optimal auction mechanism or Meyerson (1981) for a more technical treatment.

[5]By assuming a sealed-bid second-price auction, we abstract away from the ascending nature of the online auction. This turns out not to be a problem as the normal form of the ascneding auction is equivalent to that of the second-price auction in the independent private values environment.

2. $v < b_1$: Deviating affects the outcome only when player 1 raises his bid so as to win the auction at $\max\{r, b_1\}$ and receive negative surplus when he would have received zero by playing the prescribed strategy.

3. $v = b_1$: Player 1's payoffs are zero if he wins and zero if he loses regardless of his bid.

■

Given the best response of buyers, we now wish to calculate the seller's expected revenue. We begin by calculating the expected payment for a given buyer, say buyer 1. In what follows, let Y_1 denote the largest of the valuations of all players other than player 1 with $G(y)$ its distribution and associated density $g(y)$. Since v_i are independent, it follows that $G(y) = F(y)^{n-1}$. Conditional on winning the auction, buyer 1's payment is r when $Y_1 < r$ and Y_1 when $Y_1 \geq r$. Thus, his expected payment is:

$$m^c(v) = rG(r) + \int_r^v yg(y)\,dy = vG(v) - \int_r^v G(y)\,dy \tag{1}$$

Expected receipts from the seller's perspective, R^c, is just $n-$times the expectation of $m^c(v)$ for $v \geq r$. Thus, $R^c(r) = n\int_r^{\bar{v}} m^c(v) f(v)\,dv$. The seller's revenue is the net of the expected receipts, R^c, and the listing fee c. The seller then chooses r^* to solve [6]:

$$\Gamma^c = \underset{r\in[\underline{v},\bar{v}]}{Max}\{-c(r) + R^c(r)\} \tag{2}$$

Since $c(\cdot)$ is a discrete function, Γ^c is non-differentiable. To solve (2), we compare revenue from the reserve maximizing R^c with the reserve prices located along at the jump points in $c(\cdot)$. Let $\hat{r}$ denote the reserve that maximizes R^c, assumed to be unique. Direct computation yields the interior solution $\hat{r} = \frac{1-F(\hat{r})}{f(\hat{r})}$. The seller's optimal reserve r^* can be no greater than $\hat{r}$ since such a reserve would cost at least as much to list but result in a lower expected sales price. Let $k(\hat{r})$ denote the index of the reserve such that $c\left(r^{k(\hat{r})}\right) = c(\hat{r})$. It follows that:

$$r^* = \arg\max\{\Gamma^c(r^0), \Gamma^c(r^1), ..., \Gamma^c\left(r^{k(\hat{r})-1}\right), \Gamma^c(\hat{r})\} \tag{3}$$

The upshot is, if listing fees are significant enough to affect the seller's choice of reserve, she will choose a reserve at the corner of a given step. Any alternative reserve on the same step would cost her the same to list but result in a lower expected price as it would be further from $\hat{r}$. Suppose that $\hat{r} = \$14$. From Table 1 (in Appendix), the seller incurs a listing fee of 65¢. Other candidates for r^* are \$9.99 and 99¢ which result in fees of 35¢ and 25¢ respectively. However, if $\underline{v} > 99$¢, then any reserve under 99¢ is a candidate as they are all equally profitable. The seller will choose a reserve of 99¢ or less rather than $\hat{r}$ if the

[6] In the actual online environment, the seller pays a fraction λ of the final sale price to eBay. We then interpret c as the actual listing fee normalized by dividing by $1 - \lambda$.

decrease in receipts is less than the 40¢ difference in listing fee. While this may be the case for some items, these cost savings are small enough that they are unlikely to account for a 99¢ reserve in all cases. Such ambiguity is averted in the dynamic model when the seller's receipts actually increase for a reserve of 99¢ over $\hat{r}$.

It is within this framework that the presence of a bidding gap is so paradoxical. Since every buyer whose valuation exceeds a reserve of say $14 will bid their valuation, for a given realization of the random variables, if a 99¢ auction results in a sale price in excess of $14, so should the $14 reserve auction. From an empirical standpoint, for any $p > r$, the probability that the transaction price exceeds p given a reserve r can be calculated from the density of the second-order statistic of valuations to be $P\{price > p|r\} = n(n-1)\int_p^{\bar{v}} F(x)^{n-2}[1-F(x)]f(x)\,dx$, which is independent of r.

2.2 Dynamic Model with No Commitment

Lacking a credible commitment device, the seller suffers from a time-inconsistency problem. That is, if the item fails to sell in the initial auction, the optimal response has her put it on sale in the following period and reduce the reserve every subsequent period until the item eventually sells or until the reserve reaches $\underline{v}$, at which point the item will necessarily sell. To analyze this process, we consider an infinite horizon model with periods indexed $t = 1, 2, ...$ In period 1, the seller runs a second-price auction with reserve r. If the item sells, the outcome is as before. If it does not sell, the seller runs another second-price auction in period 2. In the period-2 auction, the seller must choose the reserve optimally, given the item failed to sell in period 1. In equilibrium, this reserve price will be no greater than that chosen in period 1. If the item fails to sell in period 2, the game repeats itself in period 3 and so on. As the game advances from one period to the next, all buyers discount returns accrued in the next period by factor δ_b and the seller discounts by factor δ_s, both of which are common knowledge. We simplify the problem for the time being by assuming the listing fee to be some constant c and that c is small enough that it is profitable for the seller to offer the item for auction in the initial period.

The technical treatment of this problem follows that of Fudenberg, Levine, & Tirole (1985) on sequential bargaining and Gul, Sonnenschein, & Wilson (1986) on durable goods monopoly as the three problems are closely related. In sequential bargaining, a seller who's valuation is known, makes a sequence of price offers (one per period) to a single buyer who's valuation is a random variable with distribution F. In durable goods monopoly, a producer who's (constant) marginal cost is known, makes a sequence of price offers (one per period) to a mass of consumers with unit demands where $F(x)$ represents the mass of consumers who's valuations lie below x. In both models, the strategy of a buyer involves deciding on a price at which the seller's offer is accepted. The difference is, sequential bargaining deals an uncertain environment whereas the environment of the durable goods monopoly is deterministic. The auction setting considered herein can be thought of as an $n-$buyer extension of sequential bargaining. The seller makes offers through her choice of reserve price. Buyers indicate acceptance by placing a bid in excess of the reserve. The magnitude of the bid comes into play only when more than one buyer accepts the seller's offer. In this case, the magnitude of the bid breaks ties by raising the price to the point at which only the winning bidder continues to accept.

To analyze this game, we search for symmetric *perfect-Bayesian-Nash equilibrium* strategies. We show that an equilibrium exists though the equilibrium path depends upon whether or not $\underline{v}$ exceeds c. In each case, the unique equilibrium is defined by $\sigma(H_{t-1})$, the seller's best response function, and $\beta(H_t)$, the lowest valuation type to place a bid (with positive probability) in each period t respectively given the history of of reserve prices, where $H_\tau = \{r_1, r_2, ..., r_\tau\}$. We begin by establishing essential properties of either equilibrium.

Lemma (2) establishes the Coasian nature of the game.

Lemma 2 *In any period t, there exists a minimum type $\beta_t = \beta(H_t)$, such that every buyer whose valuation exceeds β_t submits a bid.*

Proof. It is sufficient to show that if it is profitable for a bidder with valuation v to submit a bid in some period t, then bidding will also be profitable for a buyer with valuation $v' > v$. In this vein, if a buyer with valuation v submits a bid period t, then it must be that payoff from doing so exceeds the continuation payoff. Formally, this implies

$$vQ_t(v) - m_t(v) \geq \delta_b \Pi_{t+1}(v; v) \tag{4}$$

where $Q_t(z) = Q(z; H_t)$ is the probability of winning in period t given history H_t for a buyer who plays the strategy of a type-z buyer, $m_t(z)$ is the analogous expected payment for such a player, and $\Pi_{t+1}(z; v, H_t)$ is expected continuation payoff for a buyer with valuation v playing as if his type were z. Of course in equilibrium, $z = v$ as each player plays according to his type. We need to show that for every $v' > v$, we have $v'Q_t(v') - m_t(v') > \delta_b \Pi_{t+1}(v'; v')$. Because a buyer always has the option of bidding in the current period, $\delta_b[\Pi_{t+1}(v'; v') - \Pi_{t+1}(v; v)] \leq \Pi_t(v'; v') - \Pi_t(v; v)$. And since the v' type can always follow the strategy of the $v-$type player, $\Pi_t(v'; v') - \Pi_t(v; v) \leq (v' - v) Q_t(v)$. Therefore,

$$\delta_b[\Pi_{t+1}(v'; v') - \Pi_{t+1}(v; v)] \leq (v' - v) Q_t(v) \tag{5}$$

Incentive compatibility implies that in every period $v'Q_t(v') - m_t(v') \geq v'Q_t(v) - m_t(v)$ or equivalently

$$v'[Q_t(v') - Q_t(v)] - [m_t(v') - m_t(v)] \geq 0 \tag{6}$$

Adding (6) to both sides of (4) yields $v'Q_t(v') - m_t(v') - (v' - v) Q_t(v) \geq \delta_b \Pi_{t+1}(v; v)$. Rearranging terms and making use of (5) yields the desired result. ■

Analogous to the literature on durable good monopoly, β_t is the lowest valuation type to buy in period t. It must be the case that $\beta_t \geq r_t$ otherwise a player could earn negative surplus by bidding.

As a point of terminology, we refer to β_t as the posterior in period t since it refers to the highest remaining type after the period t auction has taken place. This type then becomes the common prior, u_{t+1}, in the following period from which the seller chooses her optimal reserve.

The Coasian nature of the game ultimately restricts the set of buyer types that bid in a given period. However, this should not affect a buyer's strategy conditional on submitting a bid.

Lemma 3 *Conditional upon bidding, it is a weakly dominant strategy for each bidder to bid his valuation.*

Proof. The proof is the same as that for Lemma (1). ■

This result is rather intuitive. Having already made the decision to bid in the current auction, each bidder chooses a bid to maximize surplus. By definition, a dominant strategy has each buyer playing the prescribed strategy regardless of the strategy of others, in particular if some of the others choose not to bid.

With buyers' strategies well-defined, the probability of winning the period-t auction is $Q_t(v) = G(v)$ for $v \geq \beta_t$; 0 otherwise. A buyer's expected payment is $m_t(v) = r_t G(\beta_t) + \int_{\beta_t}^{v} y g(y)\, dy$ for $v \geq \beta_t$; 0 otherwise. Since the seller's revenue in a given is just $n-$times a given buyer's expected payment, the seller chooses r_t to maximize $\Gamma(u_t) = \sum_{j=0}^{\infty} \delta^j F(u_{t-1+j})^n \left(n \int_{\beta(r_{t+j})}^{u_{t+j}} m_{t+j}(v) \frac{f(v)}{F(u_{t+j-1})^n} dv - c \right)$ where β is a known function of r in equilibrium. The following Lemma shows that this necessarily results in a downward sloping reserve-price path.

Lemma 4 *The seller's equilibrium reserve-price path is strictly decreasing.*

Proof. Suppose by way of contradiction that $r_{t+1} \geq r_t$. We will show that any player that submits a bid in $t+1$ should also bid in period t. Consider a buyer with valuation v that prefers not bid in the current auction ($v < \beta_t$). Since this buyer would have won the current auction only when no one else bids ($Y_1 < \beta_t$), not bidding implies,

$$(v - r_t) G(\beta_t) < \delta_b \Pi_{t+1}(v) \tag{7}$$

For r_{t+1} to be part of an equilibrium, it must be the case that some buyer type bids in period $t+1$. Otherwise, the seller would incur listing fee c and gain no additional revenue. She could increase expected profits by jumping ahead to the next reserve price that results in some type placing a bid. Thus, there must exist a minimum type β_{t+1} strictly below β_t that bids in period $t+1$. This type, by virtue of being the lowest type to bid, earns expected surplus

$$\Pi_{t+1}\left(\beta_{t+1}\right) = \left(\beta_{t+1} - r_{t+1}\right) G\left(\beta_{t+1}\right) \tag{8}$$

Since, as we have argued, type β_{t+1} does not bid in period t let $v = \beta_{t+1}$ in (7). Combining equations, we have

$$\left(\beta_{t+1} - r_t\right) G(\beta_t) < \left(\beta_{t+1} - r_{t+1}\right) G\left(\beta_{t+1}\right) \tag{9}$$

which is a contradiction since both sides are positive, $G\left(\beta_{t+1}\right) < G\left(\beta_t\right)$ by Lemma (2) and $r_{t+1} \geq r_t$ by hypothesis. ■

The upshot is, if the seller posts a reserve greater than the last, she won't receive any bids. Following this logic recursively, the current reserve must be lower than all previous reserve prices.

As the game goes on, if the item remains unsold the reserve may eventually be reduced to $\underline{v}$. At this point, the seller needn't reduce the reserve any farther as all buyer types will bid.

Lemma 5 *Regardless of the history, all buyer types bid when the reserve is at or below* $\underline{v}$.

Proof. We begin by asserting that the seller will never drop the reserve below her valuation of zero. At a reserve of zero, a buyer submits a bid as long as his expected surplus exceeds his continuation payoff. Since the reserve can get no lower, the continuation payoff can be no greater than the payoff of bidding. With every bidder bidding their respective valuation, each type receives $\int_{\underline{v}}^{v} G(y)\,dy \geq 0$. Thus, every buyer type bids when the reserve is zero.

We now use recursive logic to show that every type bids for any reserve less-than or equal to $\underline{v}$. Knowing that zero is the lowest possible reserve, buyers bid at reserve price $r \in (0, \underline{v}]$ as long as $vG(v) - m(v;r) \geq \delta_b \int_{\underline{v}}^{v} G(y)\,dy \Leftrightarrow [\beta(r) - r]\,G(\beta(r)) + \int_{\beta(r)}^{v} G(y)\,dy \geq \delta_b \int_{\underline{v}}^{v} G(y)\,dy$ where $\beta(r)$ is the lowest type to submit a bid. From Lemma (2), if this condition holds for the lowest buyer type β, it holds for all higher types. Thus it is sufficient to show that this condition holds for $v = \beta$, which implies,

$$[\beta - r]\,G(\beta) \geq \int_{\underline{v}}^{\beta} G(y)\,dy \tag{10}$$

where we set $v = \beta$ and $\delta_b = 1$. This condition holds for all $\beta \geq \underline{v}$ indicating all buyer types do submit bids. Furthermore, it also holds for all $r \leq \underline{v}$, establishing the desired result. ■

Lemma (5) says that the seller can induce bids from all buyer types and guarantee a sale with a reserve of $\underline{v}$ or below. It turns out that when $\underline{v} > c$, the seller will eventually reduce the reserve to $\underline{v}$ when u_t becomes sufficiently small. This guarantees the game ends with a sale. When $\underline{v} < c$, this outcome may also be possible. But it is also possible that depending upon the magnitude of c, the seller may eventually find it unprofitable to list the item for u_t sufficiently small. In such case, the last posted reserve is binding and the game ends without a sale. Regardless of the case, we know that the game ends in finite time so we can solve for the equilibrium through backward induction. This, along with our assumption that the seller's best response function σ is single-valued, guarantees our equilibrium is Markovian. That is, the seller's best response depends only on the current prior u_t and a buyer's strategy, characterized by β, depends only on the current reserve r_t.

Proposition 6 *For any $c > 0$ and any $\delta_b, \delta_s \in [0, 1]$, an equilibrium exists and it is generically unique. The equilibrium is Markovian.*

Proof. See Appendix. ■

3 Economic Implications of Dynamic Model

In this section, we describe the bidder's decision to bid in a given period and show how bidding gaps can result. We then characterize the conditions under which our model predicts a non-binding initial reserve.

3.1 Bidding Gaps

In equilibrium, the seller's reserve price path is a deterministic sequence $\{r_t\}_{t=1}^{T^*}$ with a corresponding sequence of minimum types β_t such that a buyer bids his valuation in period t if it falls in the interval $[\beta_t, \beta_{t-1})$. Since a player with valuation β_t wins the period$-t$ auction only when $Y_1 < \beta_t$ and thus pays the reserve r_t, β_t must satisfy

$$(\beta_t - r_t) G(\beta_t) = \delta_b \Pi_{t+1}(\beta_t) \tag{11}$$

where due to the descending nature of the reserve price path, the type indifferent in period t strictly prefers bidding in $t+1$. This yields a continuation payoff of:

$$\Pi_{t+1}(\beta_t) = \left(\beta_{t+1} - r_{t+1}\right) G\left(\beta_{t+1}\right) + \int_{\beta_{t+1}}^{\beta_t} G(y)\, dy \tag{12}$$

Equations (11) and (12) define a cutoff such that a buyer bids in the current period only when the payoff of doing so exceeds the opportunity cost. This makes bidding sequentially rational Furthermore, the cutoff value β_t has some interesting economic properties laid out in the following corollary.

Corollary 7 *For any reserve r along the equilibrium path, the corresponding posterior $\beta(r)$ has the following properties:*

1. $\frac{\partial \beta}{\partial r} > 0$.
2. $\beta(\underline{v}) = \underline{v}$, $\beta(r) > r \;\forall r > \underline{v}$.
3. $\beta(r; n') < \beta(r; n)$ *for* $n' > n$, $\lim\limits_{n \to \infty} \beta = r$.
4. $\frac{\partial \beta}{\partial \delta_b} > 0$, $\underset{\delta_b \to 0}{lim} \beta = r$.

Proof. See Appendix. ■

Property 1 says the minimum valuation type submitting a bid increases with the reserve price. This property is fairly trivial as it holds in the case of seller commitment where the minimum bidding type is equal to the reserve. The interesting property, 2, says that there exists a gap between the reserve price and the minimum bidding type for any reserve greater than $\underline{v}$. It is this property that perplexed the eBay sellers. Viewing the auction as a one-shot game, they seem to think that any rational bidder willing to pay $14 in a 99¢ reserve should be willing to pay $14 in a $14 reserve auction. Property 2 says this may not be the case as the minimum type that bids in a $14 reserve auction exceeds $14. Thus, we should not be surprised to see an auction with a 99¢ reserve end with a price exceeding $14 while the $14 reserve auction ends with no sale. In particular, the probability that the price exceeds $14 when the reserve is $14 is equal to the probability that the second-highest valuation exceeds $\$\beta(14)$. In contrast, when the reserve is 99¢, the probability of the price exceeding $14 is equal to the probability that the second-highest valuation exceeds $14 which is less than $\$\beta(14)$. These probabilities would be equal could the seller commit in advance.

From an empirical standpoint, property 2 says that for some $p > r$, $P\{price > p|r\} = n(n-1)\int_{\max\{p,\beta(r)\}}^{\bar{v}} F(x)^{n-2}[1-F(x)]f(x)\,dx$, which is decreasing in r for p sufficiently close to r. The upshot is, for a given reserve r, if we choose a value p just greater than r, the probability that the price exceeds p is equal to the probability that the second-highest valuation is both greater than p and large enough that it pays to bid in the current auction. Since the value β that determines whether or not it pays to bid in the current auction is both greater than r and increasing in r, by lowering the reserve by some small amount dr, we can expect to invite bids from buyers with valuations in the interval $[\beta(r-dr), \beta(r)]$, where $\beta(r-dr)$ exceeds p. By inviting more bids from buyers whose valuation exceeds p, we should see the probability that the price exceeds p increase.

Property 2 illustrates a buyer's incentive to "game" the seller by delaying bidding. Properties 3 and 4 show how the number of buyers and the level of buyer impatience curbs this incentive. As the number of buyers increases, the probability that the item goes unsold at any reserve decreases. This has the effect of reducing a buyer's incentive to delay bidding resulting a lower value of β for a given reserve. Likewise, as buyers become more impatient–indicated by a lower value of δ_b–the value of delaying bidding is reduced resulting in a lower value of β.

3.2 Non-binding Initial Reserve

With respect to the seller, the incentive to delay bidding reduces expected receipts from a given reserve in any period. To see this, consider the expected payment for a buyer in period t:

$$m_t(v) = r_t G(\beta_t) + \int_{\beta_t}^{v} y g(y)\,dy = vG(v) - (\beta_t - r_t)G(\beta_t) - \int_{\beta_t}^{v} G(y)\,dy \quad (13)$$

As in (1), we have written the expected payment as the difference between a buyer's gross and net benefit. What distinguishes (1) from (13) is the wedge $(\beta_t - r_t)G(\beta_t)$. This term is positive and does not depend on v. From (11), we see that the wedge is equal to the reservation level surplus of the indifferent bidder. Thus, any buyer that bids must be guaranteed this level of surplus, which comes right out of the seller's pocket. For this reason, at any prior u_t, the expected payment for a buyer and hence expected receipts for the seller in period t are lower than they would have been could the seller commit herself to not relist.

We want to show that the inability to commit forces the seller's initial reserve price below what it would have been under commitment. In what follows, we assume that c, the generic listing fee, lies below $\underline{v}$ so that the reserve eventually drops to $\underline{v}$. Harking back to the advice of the seller Clact in the introduction, the 99¢ reserve works best when the expected sales price is well above 99¢. Thus, we can think of our assumption as a condition guaranteeing that the expected sales price is sufficiently large. However, when $\underline{v}$ is sufficiently greater than 99¢, it is not enough to show that the reserve drops to $\underline{v}$. For the seller to want to drop the reserve to 99¢, it must be that the listing fee associated with a 99¢ reserve is below the fee paid for a reserve of $\underline{v}$. To introduce this component without affecting the existence of equilibrium, we assume that the schedule of listing fees is again a step function like that on eBay with the restriction that any reserve in $[\underline{v}, \bar{v}]$ has the same fee c. Lastly, we assume that c^0, the minimum listing fee, is strictly less than c.

The optimal reserve depends in part on the discount rate of both buyers in sellers. We saw in Lemma 7 that when buyers' discount rate is 0, they behave as if the seller could commit. Regardless of her discount rate, the seller could do no worse than to post a reserve of r^*, which maximizes revenue in a one-shot game. Thus, for the seller to start with a low reserve requires that buyers be patient enough to wait for the reserve to be reduced.

Corollary 8 *For any $c \in (0, \underline{v})$ and any $\delta_s \in [0,1]$, there exists a $\bar{\delta}$ such that if δ_b exceeds $\bar{\delta}$, $\sigma(\bar{v}) = \underline{v}$. In such case, the seller's initial reserve is no greater than $r^0 < \underline{v}$, all buyer types bid and the seller earns $\int_{\underline{v}}^{\bar{v}} M^c(v; \bar{v})\, dv - c^0$. The optimal reserve, expected selling price, and total seller revenue are lower than in the case of commitment, strictly so when $r^* > \underline{v}$.*

The cutoff value for buyers' discount factor $\bar{\delta}$ is the smallest value of δ_b such that a seller with prior $\bar{v}$ prefers to make the initial reserve non-binding when he otherwise would do so in the following period. Following from the first-order condition of the seller's two-period problem, we have:

$$\bar{\delta} = 1 - [(1 - \delta_s)\, \underline{v} + \delta_s c]\, f(\underline{v}) \tag{14}$$

The introduction included a quote from the seller Clact who claims the 99¢ auction works best for items that are expected to sell for at least \$10. In such cases, the reserve constitutes a small fraction of the item's transaction price. If $\underline{v} > r^0$, then r^0, representing a reserve of 99¢, is part of the set of optimal reserves and constitutes a small fraction of the transaction price. Thus the findings of this model are consistent with our two empirical anomalies. The only remaining question is then: Why a 99¢ reserve? Why not 1¢? It turns out that many auctions do employ a reserve of 1¢ as well, though few use any in between 1¢ and 99¢.

Given the curiousness of this result, we may be interested to consider what happens when the seller tries to post an initial reserve in excess of $\underline{v}$. A numerical example is useful here. Suppose $n = 2$, $\delta_s = \delta_b = 1$, buyer valuations are distributed uniformly on $[.2, 1.2]$, and $c < .2$ In the standard auction model, the seller chooses a reserve to maximize $n \int_r^{1.2} [2v - 1.2](v - .2)dv$, which yields $\hat{r} = .6$ and $R^c = .576$. When the seller cannot commit to a reserve in period 1, $\delta_b > \bar{\delta} = 1 - c$ implies the seller posts a non-binding reserve in the initial period. If the seller posts some reserve $r_1 > .2$ in the initial period, a buyer bids only if his valuation exceeds a cutoff, $\beta^2 (r_1)$ determined as the type indifferent between bidding with the current reserve and waiting for the following period where the seller will revert back to the optimal reserve path which calls for a nonbinding reserve. From (11) and (12), we have $\beta^2 (r_1) = 2r_1 - .2$.

Suppose the seller, in the absence of commitment, posts a period-1 reserve of $\hat{r} = .6$. Only buyers whose valuations exceed $\beta^2 (\hat{r}) = 1$ submit bids resulting in expected period-1 receipts of only $R_1 (\hat{r}) = .362$. Overall, her combined receipts from periods 1 and 2 equal $R_1 + R_2 = .533$, independent of her reserve in period 1. The figure below illustrates $R^c (r)$, $R_1 (r)$, and $R_1 + R_2$ as functions of r.. Notice that R^c is positive for all $r \in [.2, 1.2]$ and is maximized at $\hat{r} = .6$. In contrast R_1 is positive only for $r \leq .7$, since no types find it profitable to bid at any higher reserve. It is maximized at .2 with expected receipts of .533. The fact that R_1 lies below R^c for any r indicates that at any reserve, not just the optimal reserve, a seller capable of commitment does better than one that is not. $R_1 + R_2$ is constant across r at a value equal to $R_1 (.2)$ since at a reserve of .2, the item is guaranteed to sell in period 1, so the seller's period 1 receipts *are* her total receipts.

As a concluding point, we point out the commitment potential of the listing fee. Corollary (8) demonstrates the worst possible outcome for the seller. And this occurs when c is small relative to $\underline{v}$. We noted in the previous section that when $c > \underline{v}$, there exist a case where the last reserve posted is a binding reserve. In such cases, in the last period, the listing fee serves as a commitment device guaranteeing the reserve will not be reduced further. For a large enough listing fee, the seller could actually commit to $\hat{r}$–the optimal reserve in the standard model–and earn net revenues greater than under conditions in which the seller is forced to post an initial reserve of 99¢. So a seller facing patient buyers may actually prefer a larger listing fee as it confers greater pricing power or strength in bargaining.

4 Empirical Predictions

The introduction presented several alternative models that are said to explain the prevalence of the 99¢ reserve. The behavioral models of our eBay sellers claim that buyer irrationality is the cause either through a competitive impulse to win the auction or through a quasi-endowment effect. The model of Milgrom

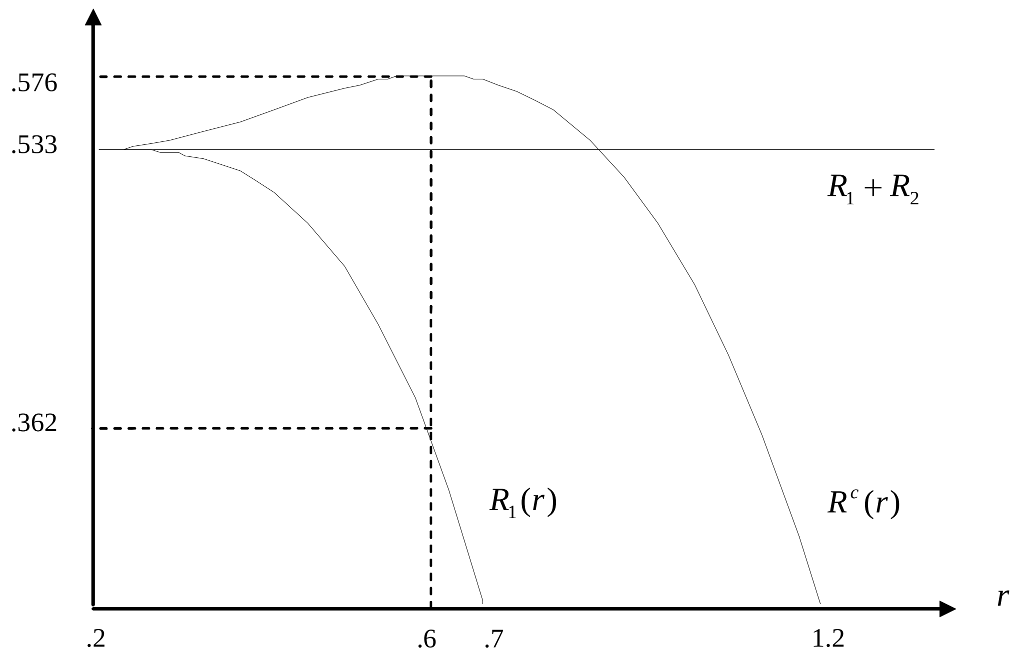

Figure 1: Comparison of $R^c(r)$ with $R_1(r)$ and $R_1 + R_2$.

Figure 1:

& Weber (1983) suggests that affiliated signals condition one bidder's willingness to pay on the bids of other players. Reiley (2004) argues that the presence of participation costs cause potential buyers to rethink their participation decision when the reserve price is high. All three of these theories result in the seller's reserve price being a small fraction of the final sale price. In what follows, we take a harder look at these theories and compare their predictions to those of the model presented herein.

We begin by distinguishing our model from the standard auction with seller commitment. In the standard auction, every buyer bids as long as his valuation exceeds the reserve r. Thus, a sale occurs as long as X_1, the first-order statistic of buyer valuations, exceeds r. Conditional on a sale occurring, the sales price is equal to r for all realizations of X_2, the second-order statistic, less than r. For all realizations of X_2 greater than r, the sales price is equal to X_2. Thus, the density of sales prices is identical to the density of X_2 for values greater than r with a mass point at r equal to the probability of $X_2 < r$ conditional on $X_1 > r$.

In contrast to the standard model, in our model, potential buyers bid if their valuations exceed some cutoff β which exceeds the reserve for all reserves greater than $\underline{v}$. Thus a sale takes place only if X_1 exceeds β. However, the sales price can take on values less than β when $X_2 < \beta$, in which case sales price is r. As before, when X_2 exceeds β, the sales price is equal to X_2. Thus, the density of sales prices is identical to the density of sales prices in the standard model for values of X_2 greater than β with a hole over the interval (r, β) and a mass point at r equal to the mass point in the standard model plus the area of the hole. The presence of this hole indicates the presence of positive selection. If we compare a 99¢ reserve auction to a \$14 reserve auction and restrict attention to outcomes in which the sales price exceeds \$14, the average sales price will be higher for the \$14 auction since the calculation only includes values of X_2 exceeding β. This positive selection gives rise to a bidding gap. As explained in the introduction, this is tendency for a 99¢ auction to result in a price exceeding \$14 while a \$14-reserve auction results in no sale at all. This happens when $14 < X_2 < X_1 < \beta(14)$. To empirically test for the presence of a bidding gap, we consider $\Omega(p; r) \equiv P\{price > p|r\}$ for $p > r$. In our model $\Omega(p; r) = n(n-1)\int_{\max\{p,\beta(r)\}}^{\bar{v}} F(x)^{n-2}[1-F(x)]f(x)\,dx$ is decreasing in r for values of p near r while in the standard model $\Omega(x; r) = n(n-1)\int_{p}^{\bar{v}} F(x)^{n-2}[1-F(x)]f(x)\,dx$ is independent of r.

4.1 Participation costs

The model of Levin & Smith (1994) is identical to the standard auction model only with some exogenous participation cost, incurred prior to each player learning his valuation. Nature (or some exogenous process), selects N the number of potential buyers. The participation cost, e, is assumed to be large enough that if all buyers participate, each receive negative expected surplus. Since buyers do not know their valuation before incurring the participation cost, a buyer will choose not to do so if he thinks all others will participate. But if they all thought that way, a buyer could deviate and profit from being the only bidder in the auction. The symmetric equilibrium calls for each buyer to participate with probability q, where q satisfies an ex-ante zero-surplus condition. The number of participants, n, is then a binomial random variable with mean qN. Since each buyer receives zero expected surplus in equilibrium, net social surplus is equal to the seller's revenue. With revenue maximization implying net social surplus maximization, the seller will not choose a reserve price that precludes any surplus-increasing transaction. For a seller with a valuation of zero, this implies that a binding reserve is never optimal.

With respect to bidding gaps, the effect of the reserve on Ω is felt only through its effect on participation. This is because an ex-ante participation cost does not screen participants so selection is random. However, it turns out that the effect of r on Ω is ambiguous. To see this, notice that $\Omega(p;r) = \sum_{n=0}^{N} P\{X_{2:n} > p|n\} P\{N=n\}$, where the first term is the probability of the second-order statistic of n bidders exceeds p and the second term is the probability of there being n bidders. As mentioned earlier, the effect of r is felt only with respect to the second term, where $P\{N=n\} = \binom{N}{n} q^n (1-q)^{N-n}$. We have that $\frac{\partial P\{N=n\}}{r} = \binom{N}{n} q^{n-1} (1-q)^{N-n} [n - Nq] \frac{\partial q}{\partial r}$, where $\frac{\partial q}{\partial r}$ is negative and $[n - Nq]$ is negative over small values of n and positive over large. Thus $\frac{\partial \Omega}{\partial r}$ is a weighted sum of $\frac{\partial P\{N=n\}}{\partial r}$ which is generally ambiguous. To test for the presence of such participation costs, we simply regress Ω on r while controlling for participation. The model of Levin & Smith predicts no relationship once participation is controlled for while our model does as participation has already been controlled for.

On the other hand, if the cost of participation is incurred after buyers know their valuations, only those with high valuations will participate. This gives rise to the same type of positive selection as in our model. In a result originally due to Riley & Samuelson, there is a duality between a reserve price and an ex-post entry fee. If the seller in a standard auction wishes to exclude all buyers with valuations below r^*, the seller can post a reserve of r^* or charge an entry fee $e(r^*)$ satisfying $e(r^*) = \int_{\underline{v}}^{r^*} G(y)\, dy$. This entry fee extracts all surplus from a buyer with valuation r^*, insuring that only those with valuations above r^* participate. But in the environment of the online auction, the participation cost is exogenous–thought to be the time cost of bidding–and is sunk. Consider the seller's problem in a standard auction where buyers face an ex-post participation cost e. Before the seller can even think about posting a reserve, all buyers with valuations below e are already excluded. If e is greater than $\int_{\underline{v}}^{r^*} G(y)\, dy$ the participation cost has already excluded a larger set of buyers than the seller would have chosen herself. Her best response would be to post a non-binding reserve. If e is less than $\int_{\underline{v}}^{r^*} G(y)\, dy$, the seller can exclude all types up to r^*

with a reserve $r(e)$ satisfying:

$$e = \int_{r(e)}^{r^*} G(y)\,dy \tag{15}$$

Notice that (15) defines a gap between the reserve price, $r(e)$, and the lowest type submitting a bid, r^*. This leads to a hole in the density of sales prices just as in our model. One way in which we can distinguish between the two models is to look at auctions of varying length. On eBay, the length of the auction is predetermined by the seller as either 3, 5, or 7 days. If buyers incur the participation cost after discovering their valuation, then the cost is in the monitoring of the auction and the placing of bids. It then stands to reason that longer auctions result in greater costs. In this case, the effect on overall participation is ambiguous as greater length allows more buyers to find out about the auction but greater costs deter more people. However, greater costs lead to a greater selection effect. Thus, if ex-post participation costs are present, Ω should be increasing in auction length, controlling for participation. In our model, the length of the auction does not effect Ω when controlling for participation.

4.2 Affiliated values

In the affiliated values story, bids placed by other players serve as signals, positive or neutral, of the item's true value. The greater the number of bids placed prior to the end of the auction, the greater a buyer's willingness to pay will escalate. Thus to encourage more bidding, the seller should start the auction at a low enough reserve. Affiliated values can account for bidding gaps in the event that a player's ex-ante expected valuation is less than the level of a reserve, but becomes greater than that level conditional upon observing the bid of other players. Though commitment may still play a role in a model of affiliated values, we may rule out affiliated values as a cause of bidding gaps by restricting attention to goods who's value does not depend on the perception of others or that do not tend to have resale value.

4.3 Quasi-endowment & competition effects

Since irrationality can be said to explain any counter-intuitive behavior, we must be rigorous in evaluating the explanatory power of the behavioral hypotheses posed by the eBay sellers. The idea that bidders respond to competition from other bidders Heyman et. al.(2004) dub an "opponent effect." To test for the presence of an opponent effect involves looking at bidding behavior as the number of competing bidders varies. Heyman et al hypothesize that as the number of competing bidders increases, the opponent effect should be greatest and thus result in greater response by an individual bidder. Of course, our model also predicts a higher price the greater the number of participants. We can distinguish the two models by looking at bid increments, the difference between a bidder's second-to-last and last bid and the price he ultimately pays. In the presence of an opponent effect, there should be a positive relationship between bid increment and the number of bidders. In our model, there should be no relationship.

The second behavioral hypothesis which claims that bidders become attached to an item during the course of the auction is dubbed a "quasi-endowment effect" by Heyman et. al. The well known concept of the "endowment effect"[7] says that ownership of a particular item raises the minimum amount the owner will accept to sell the good above what he would have initially paid to buy it. During the course of an auction, the leading bidder may experience a sense of ownership–quasi-ownership–over the good up for bid. Thus when another bidder comes along and tops his bid, he is compelled to increase his bid due to the sense of loss he would experience from not obtaining the good. Heyman et al. hypothesize that quasi-endowment is greatest the longer a participant is involved in an auction and in particular the longer he/she is actually the leading bidder. Thus the more time spent as the leading bidder, the greater the amount such a bidder should be willing to bid once overtaken by another bidder. In an ascending auction version of our model, we would presume such bid increases to be random and thus independent of the time spent in the lead. The test between our model and a model of quasi-endowment effects would simply check for a relationship between time in the lead and final bid increment.

5 Conclusion

...

6 Appendix

6.1 Table 1

Reserve Price	Listing Fee
\$0.01-\$0.99	\$0.25
\$1.00-\$9.99	\$0.35
\$10.00-\$24.99	\$0.65
\$25.00-\$49.99	\$1.20
\$50.00-\$199.99	\$2.40
\$200.00-\$499.99	\$3.60
\$500.00 or more	\$4.80

6.2 Proof of Proposition 6

The result must be proven for two distinct cases. We begin with the case of $\underline{v} > c$. To prove existence by backward induction, we first establish the existence of a finite end time.

[7] Kahneman, Knetsch, & Thaler (1990).

Lemma 9 *When $\underline{v} > c$, there exists a T^L such that if the item fails to sell in the first T^L periods, the seller drops the reserve below $\underline{v}$ in the following period.*

Proof. We first show that when u_t is sufficiently small, the seller posts a non-binding reserve regardless of the prior history. We do this by demonstrating that when buyers play myopically ($\delta_B = 0$), there exists a u^* such that for any $u_t \leq u^*$, the optimal reserve is no greater than $\underline{v}$. This argument is sufficient for our purposes as the incentive to drop the reserve is greater when buyers are not myopic. Myopic buyers bid whenever their valuation exceeds the reserve. Thus the current reserve becomes the posterior if the item goes unsold. Starting with a posterior of u_t, the seller's expected revenue is bounded by the following:

$$B(r) = n \int_r^{u_t} [vf(v) + F(v) - F(u_t)] F(v)^{n-1} dv + F(r)^n \delta_S (r - c) \quad (16)$$

The first-order condition is negative if $F(u_t) < [(1-\delta_s)\underline{v} + \delta_s c] f(\underline{v})$, which is true for sufficiently small u_t as $f(\underline{v})$ is bounded away from zero. This establishes the existence of u^*.

Next we show that the posterior gets to below u^* in finite time. Once again, considering myopic buyers puts a lower bound on the speed of descent of the posterior. The first-order condition establishes the distance between the prior and the chosen reserve: $[F(u_t) - F(r)] > [(1-\delta_s)\underline{v} + \delta_s c] \underline{f}$, where $\underline{f}$ is the minimum of f. Thus the distance between the posterior and the chosen reserve is bounded from below whenever the optimal reserve is interior. When it is not interior, then by definition u_t is below u^*. ■

Given the result of Lemma (**??**), we can solve for the equilibrium through backward induction. We begin by establishing cutoff values for the seller's prior such that if say $u_t < u^k$, the seller chooses a reserve price that ends the game in at most $k-1$ additional periods.

For readability and ease of notation, let $n \int_{\beta_t}^{u_t} M^c(v; u_t)\, dv$ denote the expected period-t receipts of a seller in the standard model when u_t is the prior. $M^c(v; u_t) = [vf(v) + F(v) - F(u_t)] G(v)$. In the absence of commitment, in any period t, the seller's expected receipts are $R_t = n \int_{\beta_t}^{u_t} [M^c(v; u_t) - (\beta_t - r_t) G(\beta_t) f(v)]\, dv$.

Let $\beta^2(r)$ denote the type indifferent between bidding today and bidding in the following period when the following period's reserve will be $\underline{v}$. $\beta^2(r)$ is unique and increasing in r. Let $\Gamma^2(u)$ denote the seller's unconditional expected payoff with prior u when constrained to post a reserve of $\underline{v}$ in the following period. That is, when the lowest type to bid is $\beta^2(r)$. Suppose $\sigma^2(u)$ is the argmax of $\Gamma^2(u)$, which we assume to be unique. Since σ^2 is increasing in u, we can uniquely define u^2 such that $u^2 = \max\left\{u < \bar{v} | \Gamma^2(u) = n \int_{\underline{v}}^{u} M^c(v; u)\, dv - c\right\}$. u^2 is the highest prior such that the seller prefers to reduce the reserve to $\underline{v}$ in the current period when constrained to do so in the following period. Thus for any $u_t < u^2$, the game will last exactly one more period.

Now consider some generic $k > 2$. Let $\beta^k(r)$ be the type indifferent between bidding today with reserve r and in the following period with reserve

$\sigma^{k-1}\left(\beta^k(r)\right)$. If σ^{k-1} and $\beta^{k-1}$are increasing functions, then $\beta^k(r)$ is well-defined and unique. Consider $\Gamma^k(u) = -c+\max\left\{n\int_{\beta^k(r)}^{u}\left[M^c(v;u) - \left[\beta^k(r) - r\right]G\left(\beta^k(r)\right)f(v)\right]dv + \delta_s\Gamma^{k-}\right.$ Let $\sigma^k(u)$ be the argmax of $\Gamma^k(u)$. Since σ^k is an increasing function of u, u^k is uniquely defined such that $u^k = \max\left\{u \leq \bar{v}|\Gamma^k(u) = \Gamma^{k-1}(u)\right\}$. u^k is then the highest prior such that the seller prefers to reduce the reserve to some value that has the reserve drop to $\underline{v}$ in $k-1$ periods. By induction, we can show that for any $k > 2$, if β^{k-1} and σ^{k-1} are unique and increasing, then σ^k is increasing in u, which implies u^k is unique.

Lemma (**??**) proved the existence of u^* such that for $u < u^*$, the seller drops the reserve to $\underline{v}$ regardless of the history. We prove the rest by backward induction on u.

Lemma 10 *If $u_t \in [\underline{v}, u^2)$, then $\sigma_t(H_{t-1}) = \underline{v}$.*

Proof. Choose ε_1 sufficiently small such that for any $\beta \in \left[u^*, u^2\right]$, $n\int_\beta^{\beta+\varepsilon_1} M^c(v;\beta+\varepsilon_1)\,dv + \delta_s\left[n\int_{\underline{v}}^{\beta} M^c(v;\beta)f(v)\,dv - F(\beta)^n c\right] < n\int_{\underline{v}}^{\beta+\varepsilon_1} M^c(v;\beta+\varepsilon_1)\,dv$ and $u^*+\varepsilon_1 < u^2$. This says that for ε_1 small enough, an upper bound on two-period revenue resulting from a non-binding reserve is smaller than that of a reserve of $\underline{v}$. We claim that at time t, for some history, $u_t \in (u^*, u^*+\varepsilon_1]$ implies a best response of $\underline{v}$. She can always guarantee herself $\int_{\underline{v}}^{u_t} M^c(v;u_t)\,dv$.

Assume she posts a reserve of $r_t > \underline{v}$ such that β_t falls in the interval $(\underline{v}, u^*]$. Then $r_{t+1} = \underline{v}$ and a buyer bids in period t if his valuation exceeds $\beta^2(r_t)$. Since the seller will drop the reserve to $\underline{v}$ in the following period and $u_t \leq u^2$, she does so in the current period by definition of u^2. Alternatively, posting a reserve leading to $\beta_t \in (u^*, u^*+\varepsilon_1]$, she obtains at most $n\int_{u^*}^{u^*+\varepsilon_1} M^c(v;u^*+\varepsilon_1)\,dv + \delta_s\left[n\int_{\underline{v}}^{u^*} M^c(v;u^*)\,dv - F(u^*)^n c\right] < n\int_{\underline{v}}^{u^*+\varepsilon_1} M^c(v;u^*+\varepsilon_1)\,dv$ by construction. Therefore, for any history at time t such that $u_t \in (\underline{v}, u^*+\varepsilon_1]$, $r_t = \underline{v}$, $\beta_t = \beta^2$, and $\Gamma(u_t) = \Gamma^1(u_t)$. Given this result, we can prove by induction that this is true for any $u_t \in [\underline{v}, u^2)$. ■

Lemma 11 *If $u_t \in [u^2, u^3)$, then $\sigma_t(H_{t-1}) = \sigma^2(u_t)$.*

Proof. The proof is analogous to that of Lemma (10). That in mind, let $u_t \in [u^2, u^3)$ and define ε_2 small enough that for every $\beta \in [u^2, u^3)$, $n\int_\beta^{\beta+\varepsilon_2} M^c(v;\beta+\varepsilon_2)\,dv + \delta_s[\Gamma^2(\beta) - F(\beta)^n c] < \Gamma^2(\beta+\varepsilon_2)$ and $u^2+\varepsilon_2 < u^3$. We claim that if at time t, for some history, $u_t \in [u^2, u^2+\varepsilon_2)$ implies a best response of $\sigma^2(u_t)$. The seller can guarantee herself $\Gamma^2(u_t)$ as buyers with valuation $\beta^2(r_t)$ bid in the current period since they expect a reserve of $\underline{v}$ in period $t+1$.

Posting a reserve that leads to $\beta_t < u^2$, the seller earns $\Gamma^2(u_t)$ since only buyers with valuations exceeding $\beta^2(r_t)$ bid in the current period anticipating a reserve of $\underline{v}$ in the following period. Therefore, an upper bound on what she gets when $\beta_t \geq u^2$, is $n\int_{u^2}^{u_t} M^c(v;u_t)\,dv + \delta_s[\Gamma^2(u_t) - F\left(u^2\right)^n c] < \Gamma^2(u_t)$. So the seller will never make such an offer. Thus, when $u_t \in [u^2, u^2+\varepsilon_2)$, $r_t = \sigma^2(u_t)$,

$\beta_t = \beta^2(r_t)$ and $\Gamma(u_t) = \Gamma^2(u_t)$. Given this result, we can prove by induction that this is true for any $u_t \in [u^2, u^3)$. ■

The remainder of the proof involves showing for any $k > 2$, $u_t \in [u^k, u^{k+1})$, $\sigma_t(H_t) = \sigma^k(u_t)$. As before, choose ε_k small enough such that for every $\beta \in [u^k, u^{k+1})$, $n\int_{\beta}^{\beta+\varepsilon_k} M^c(v; \beta+\varepsilon_k) f(v)\, dv + \delta_s \left[\Gamma^k(\beta) - F(\beta)^n c\right] < \Gamma^3(\beta+\varepsilon_k)$. To show that the seller can always guarantee herself $\Gamma^k(u_t)$, we need to show that when the seller posts a reserve of $r_t = \sigma^k(u_t)$, a buyer with valuation $\beta^k(r_t)$ bids. If not, then $u_{t+1} > \beta^k(r_t)$ which implies $r_{t+1} > \sigma^{k-1}(u_{t+1}) > \sigma^{k-1}\left(\beta^k(r_t)\right)$. But then a buyer of type $\beta^k(r_t)$ would not bid in the following period since by definition, such a buyer is indifferent between bidding at r_t and at $\sigma^{k-1}\left(\beta^k(r_t)\right)$. When the reserve in the following period $r_{t+1} > \sigma^{k-1}\left(\beta^k(r_t)\right)$ he prefers not to bid. But this contradicts the definition of β^k, thus he bids at the current reserve r_t, so the seller earns $\Gamma^k(u_t)$. We complete the proof using induction on k for at most T^L periods. And since, by construction, r_t depends only on the prior u_t, the equilibrium is Markovian.

We move now to the case where $\underline{v} \leq c$. This encompasses the "no gap" case ($\underline{v} = 0$) considered in FLT and in the durable goods monopoly literature. In those papers, the analysis was complicated by the fact that the game need not end in finite time. Due to the presence of a listing fee in our analysis, the game still necessarily ends in finite time so we may again solve for the equilibrium path via backward induction.

Lemma 12 *When $\underline{v} \leq c$, there exists a T^H such that the game ends in at most $T^H + 1$ periods.*

Proof. There are two cases two consider. One, when the seller drops the reserve to $\underline{v}$ and the game ends with a sale. In the other, the seller stops listing the item for sufficiently small priors so that the game ends with a binding reserve and may not result in a sale. Notice that when $u_t = \underline{v}$, the seller's revenue is at most $\underline{v} - c < 0$. Thus there must exist a u^1 such that whenever $u_t < u^1$, the seller prefers not to list the item and $T^H = t - 1$. Let $\hat{r}(u_t)$ denote the optimal reserve given prior u_t in a one-period game. It follows that $u^1 < \bar{v}$ satisfies $\int_{\hat{r}(u^1)}^{u^1} M^c(v; u^1)\, dv - c = 0$.

As in the case of $\underline{v} > c$, define u^* such that for $u_t \leq u^*$, $\sigma(u_t) = \underline{v}$. $u^* = \max\left\{u_t \leq \bar{v} | \frac{\partial \Gamma(u_t)}{\partial r_t} < 0\right\}$. There is no natural comparison between u^* and u^1 as one depends on the slope of Γ while the other on its magnitude. If $u^1 > u^*$ then for any $u_t < u^1$, the seller does not list the item and the game ends. If $u^* > u^1$, then for any $u_t \in [u^1, u^*]$, $\hat{r}(u_t) = \underline{v}$ and the game ends in one more period while if $u_t < u^1$ the game ends immediately.

We need now only show that the game reaches $\max\left\{u^*, u^1\right\}$ in finite time. When $u^* > u^1$, the proof is the same as before: when r_t is an interior solution, the rate of descent is bounded below, when it is not, $u_t < u^*$ and we're done. When $u^1 > u^*$, if u^* is reached in finite time, then so is u^1. ■

Within this subcase, there are two distinct equilibria depending upon which is larger u^* or u^1. If $u^1 > u^*$, then there exists a u^2 such that for any $u_\tau \in [u^1, u^2)$, $\sigma(u_\tau) = \hat{r}(u_\tau) < u^1$. We can then work backward to show existence and uniqueness of u^k, $k > 2$, such that if $u_t \in [u^k, u^{k+1})$, the seller chooses a reserve that has the game end in at most k periods. If $u^* > u^1$, then there exists a distinct u^2 such that for any $u_\tau \in [u^1, u^2)$, $\sigma(u_\tau) = \underline{v}$. The rest of the proof in either distinct case is analogous to the proof for when $\underline{v} > c$.

6.3 Proof of Corollary 7

1. From (11) and (12), differentiating with respect to r yields:

$$\frac{\partial \beta_t}{\partial r_t} = G(\beta_t) / [(1-\delta_b) G(\beta_t) + (\beta_t - r_t) g(\beta_t)] > 0 \tag{17}$$

2. When $r_t = \underline{v}$, Π_{t+1} is necessarily zero. This requires that the left-hand side of (11)is also zero, which implies $\beta_t = \underline{v}$. When $r_t > \underline{v}$, $\Pi_{t+1} > 0$, which requires the left-hand side of (11) to be positive, which implies $\beta_t > r_t$.

3. $\frac{\Pi_{t+1}(v)}{G(v)} = E\left[v - r_{t+1} I\left\{Y_1 < \beta_{t+1}\right\} - Y_1 I\left\{Y_1 > \beta_{t+1}\right\} | Y_1 < v\right]$ is decreasing in n due to stochastic dominance, where I represents the indicator function. It then follows from (8), that $\beta_t - r_t$ is smaller for larger values of n.

4. From (11) and (12), differentiating with respect to δ_b yields:

$$\frac{\partial \beta_t}{\partial \delta_b} = \Pi_{t+1}(\beta_t) / [(1-\delta_b) G(\beta_t) + (\beta_t - r_t) g(\beta_t)] > 0 \tag{18}$$

6.4 Proof of Corollary 8

We need to show that for δ_b sufficiently large, the seller's best response at the initial node, $\sigma(\varnothing)$, is $\underline{v}$. Since in the final period T^L, the minimum type is $\underline{v}$, we consider the second-to-last period. Posting a reserve r, β satisfies $(\beta - r) G(\beta) = \delta_b \int_{\underline{v}}^{\beta} G(y)\, dy$. The seller then chooses $r \geq \underline{v}$ (and consequently $\beta \geq \underline{v}$) to maximize $\int_\beta^u M^c(v;u)\, dv - \delta_b \int_{\underline{v}}^{\beta} G(y)\, dy + \delta_s \left[-F(\beta)^n c + \int_{\underline{v}}^{\beta} M^c(v;u)\, dv\right]$. $\sigma(u) = \underline{v}$ when the first-order condition, evaluated at $\underline{v}$ is negative, meaning $F(u) < [(1-\delta_s)\underline{v} + c] f(\underline{v}) + \delta_b$. We want to show that when δ_b is sufficiently large, this condition holds for any $u \leq \bar{v}$. For $u = \bar{v}$, this is equivalent to $\delta_b > 1 - [(1-\delta_s)\underline{v} + \delta_s c] f(\underline{v}) \equiv \bar{\delta}$. We need only now show that $\bar{\delta} \leq 1$, which is true as long as $[(1-\delta_s)\underline{v} + \delta_s c] f(\underline{v}) \geq 0$ which is necessarily true.

References

[1] Bajari, Patrick; Hortacsu, Ali. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics*, Summer 2003, Vol. 34 No. 2, p329.

[2] Coase, R.H. “Durability and Monopoly.” *Journal of Law and Economics*, April 1972, Vol. 15 No. 1, p143.

[3] Fudenberg, D.; Levine, D. and Tirole, J. “Infinite-horizon Models of Bargaining with One-sided Incomplete Information,” in Alvin E. Roth ed., *Game Theoretic Models of Bargaining*. Cambridge: Cambridge Press, 1985, pp. 73-98.

[4] Gul, F.; Sonnenschein, H. and Wilson, R. "Dynamic Monopoly." *Journal of Economic Theory*, June 1986, Vol. 39 No. 1.

[5] Heyman, J.; Orhun Y; Ariely D. "Auction Fever: The Effect of Opponents and Quasi-Endowment on Product Valuations." *Journal of Interactive Marketing,* Autumn 2004, vol 18 no 4.

[6] Kahneman, D.; Knetsch, J.L.; Thaler, R.H. "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy,* vol 98 no 6, pp. 1325.

[7] Krishna, V. "Auction Theory." San Diego, Ca: Academic Press, 2002.

[8] Levin, D; Smith, J. "Equilibrium in Auctions with Entry." *American Economic Review*, 1994. Vol. 84 No. 3.

[9] Meyerson, Roger B. "Optimal Auction Design." *Mathematics of Operations Research*, February 1981, Vol 6 No 1, p 58.

[10] Milgrom, P; Weber, R. "A Theory of Auctions and Competitive Bidding." *Econometrica*, September 1982. Vol. 50 No. 5. p 1089.

[11] Reiley, David., “Experimental Evidence on the Endogenous Entry of Bidders in Internet Auctions,” U. of Arizona working paper.

[12] Riley, John G.; Samuelson, William F. “Optimal Auctions.” *The American Economic Review*, June 1981, Vol. 71 No. 3, p381.

[13] Schelling, Thomas. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press, 1960.

[14] Vickrey, William. “Counterspectulation, Auctions, and Competitive Sealed Tenders.” *The Journal of Finance*, March 1961, Vol. 16 No. 1, p8.